Jitter Studies for a 2.4 GeV Light Source Accelerator Using LiTrack

G. Penn

Center for Beam Physics, Lawrence Berkeley National Laboratory, Berkeley, CA 94720
(Dated: March 19, 2010)

Electron beam quality is an important factor in the performance of a free electron laser (FEL). Parameters of particular interest are the electron beam energy, slice emittance and energy spread, peak current, and energy chirp. Jitter in average energy is typically many times the slice energy spread. A seeded FEL is sensitive not only to these local properties but also to factors such as shot-to-shot consistency and the uniformity of the energy and current profiles across the bunch. The timing and bunch length jitter should be controlled to maximize the interval of time over which the electron beam can be reliably seeded by a laser to produce good output in the FEL. LiTrack, a one-dimensional tracking code which includes the effect of longitudinal wakefields, is used to study the sensitivity of the accelerator portion of a 2.4 GeV FEL to sources of variability such as the radio frequency (RF) cavities, chicanes, and the timing and efficiency of electron production at the photocathode. The main contributors to jitter in the resulting electron beam are identified and quantified for various figures of merit.

I. INTRODUCTION

As part of the design process for an FEL, the quality and repeatability of the electron beam is an important aspect to study. Electron beam quality directly affects the FEL performance, and for a seeded FEL in particular the shot-to-shot consistency imposes tight tolerances on sources of jitter. Uniformity in the energy and current profiles are desirable properties as well. A model beamline which is a candidate for a soft X-ray FEL at moderate electron energy [1] is analyzed for the impact of photocathode, radio frequency (RF) cavities, and chicane errors on beam quality. The sensitivity of the electron beam to fluctuations in the beamline is analyzed using LiTrack [2]. LiTrack is a one-dimensional tracking code which runs in Matlab [3]. LiTrack tracks particle energy and longitudinal position, models basic accelerator components, and allows for prescribed longitudinal wakefields. An additional feature of modeling jitter in chicanes was implemented by Paolo Craievich at Elettra. The physics model is most appropriate for highly relativistic beams. There is no modeling of space charge or CSR, so some physics is left out that might contribute to sensitivities to linac parameters.

Although some physics processes which are important at low energies are left out of these simulations, most of these processes will affect transverse matching and should have little coupling to the longitudinal profile. The main parameter which affects the response of the beam to self-forces as it leaves the photocathode is the accelerating gradient, so the low-energy portion of the lattice should be checked for additional sensitivities to the electric field at the cathode. This study does offer a useful, generic look at the interactions between errors in different elements of the beamline as the beam variations evolve in longitudinal phase space.

To add to this generic analysis, later sections add spatial and temporal correlations to the beamline errors. We consider a pessimistic case where most RF cavities are grouped into 9-cavity modules with identical errors, thus reducing the statistical smoothing from having independent errors. Such correlations could result if each module was powered from a single klystron, rather than having an independent power source for each RF cavity. We also consider the possibility of CW operation at a high repetition rate, facilitated through the use of superconducting RF for acceleration. Under such a mode of operation, and given the high Q of superconducting RF cavities, the spectrum of fluctuations should be dominated by acoustic frequencies. At these frequencies, the time evolution of errors would be resolved and could even be controlled through conventional feedback loops. We conclude the analysis by considering a typical power spectrum for the time evolution of each type of beamline error, and applying a model for a generic feedback system having a time resolution of either 1 ms or 0.1 ms. The charge per bunch jitter alone has been modeled in a more pessimistic way as having no shot-to-shot correlations, because the physical basis for such fluctuations is different and may depend on details such as cathode material and the photocathode laser design. While many other jitter studies of linacs using superconducting RF have been performed, most of these use pulsed operation [4, 5], and even for those that can be run CW [6] the analysis is similar to that for normal conducting RF; from the literature, it appears that the temporal evolution of errors in CW linacs and the use of rapid feedback systems have not been considered previously.

II. SIMULATION TECHNIQUE

To analyze the sensitivity of the electron beam to errors, the entire beamline must be considered. The beamline is modeled in LiTrack from the photocathode through the spreader, based on a preliminary injector design provided by Steven Lidia and on a reference linac design for a 2.4 GeV next-generation FEL. A total of 100,000 macroparticles are tracked in each LiTrack runs.

Special considerations must be made for the fact that LiTrack does not adequately model low-energy dynamics, where velocities are significantly below the speed of light, space charge becomes significant, and mixing can occur between longitudinal and transverse dynamics. Slight perturbations to a few linac parameters of the reference beamline were made; also, rather than starting from the expected beam at the cathode, we begin with a reasonable distribution at 40.8 MeV energy and run it backwards through the injector using LiTrack. This yields a mock distribution that replaces the expected initial electron beam leaving the cathode. This procedure is necessary to simulate the entire injector and linac within the physics model used by LiTrack and still reproduce the desired beam profile at the end of the accelerator. We note that the reference injector configuration is designed to yield the distribution at 40.8 MeV that we use here. The fictional initial distribution will match into the final distribution even though the physics modeling at low energies is inaccurate. While this method requires the use of an unphysical beam at the cathode with an unusual distribution, it should suffice for the purpose of pursuing jitter studies and setting tolerances to sources of jitter. Note that any interactions between space charge and longitudinal dynamics will be neglected by this procedure.

A more essential concern is the lack of non-relativistic velocities, especially because the injector design envisions using 650-MHz RF bunching cavities to shorten the beam through velocity compression. As a simplified model for motion which is not ultra-relativistic, the drift sections between the cathode and buncher, and between the buncher and first acceleration cavity, are implemented as generalized chicanes whose $R_{56} = -L\gamma_0(\gamma_0^2 - 1)^{-3/2}$, where L is the drift length and γ_0 is the reference energy (here fairly close to unity). Two nonlinear terms are kept as well, with $T_{566}/R_{56} = -(3/2)\gamma_0^2/(\gamma_0^2 - 1)$, and $U_{5666}/R_{56} = (1/2)\gamma_0^2(4\gamma_0^2 + 1)/(\gamma_0^2 - 1)^2$. Note that in the ultra-relativistic limit this gives $T_{566}/R_{56} = -3/2$ and $U_{5666}/R_{56} = 2$, the same ratios as typically hold for chicanes. However, the value of R_{56} scales as $1/\gamma_0^2$ which makes this effect unimportant in drifts at large energies.

Many sources of jitter are examined in the LiTrack simulations. Two fluctuating parameters of the initial beam are the charge per bunch and the time of impact of the photocathode laser on the photocathode. The error in this time of impact, as well as the final timing of the electron bunch, is measured with respect to a fixed, ideal reference time that governs the entire system, including RF. The phase space distribution of the electrons leaving the cathode is fixed. Shifts in these parameters are particularly important because they continue to affect beam propagation throughout the injector and linac, altering wakefields and the RF phases experienced by the beam. All RF cavities are allowed to fluctuate in both peak accelerating gradient and phase. In addition, the strength of the chicanes in the bunch compressor and the spreader fluctuate. While low-energy drift sections are modeled as generalized dispersive sections, their parameters are fixed by the drift length and do not have jitter. We first consider errors that are completely uncorrelated, both from shot-to-shot and among different beamline elements. We then model errors as evolving over time and being correlated in space, and also incorporate a model for generic feedback systems with different response times. Spatial correlations are implemented by assuming that RF cavities in linac 2 are grouped into cryomodules of 9 cavities each, and that all RF cavities in a single cryomodule have identical power and phase errors, although neighboring cryomodules are still uncorrelated. Similarly, all third harmonic RF cavities are taken to have identical errors. This increases fluctuations in the final beam by reducing the number of statistically independent errors. The temporal evolution of errors is modelled using power spectral density (PSD) functions in a frequency range from 1 to 300 Hz, grouped into 1 Hz bins.

III. CONFIGURATION AND NOMINAL PARAMETERS

The beamline design uses a nominal charge per bunch of 0.8 nC. The source of electrons is a photocathode located in a 100 MHz RF cavity which yields a maximum energy gain of 0.753 MeV. After this, the remainder of the injector section consists of a drift section, a buncher RF at 650 MHz for velocity compression, a second drift, and the injector acceleration stage which consists of 6 RF cavities at 1300 MHz and accelerates the beam up to 40.8 MeV. After the injector, the beam passes through linac 1 which has 18 RF cavities at 1300 MHz and accelerates the beam up to 240 MeV, then the third harmonic RF at 3900 MHz which has 7 cavities, the chicane for bunch compression, linac 2 which has 162 RF cavities and accelerates the beam to the final energy of 2.4 GeV, and the spreader which has nonlinear dispersion but $R_{56} = 0$.

The bunching section consists of a drift section after the cathode that is taken to have a length of 1 m, a 650 MHz cavity that yields a maximum energy gain of 0.337 MeV and whose phase is shifted -110 degrees, and a second drift section after the buncher cavity which is taken to have a length of 0.85 m. The second drift section is more crucial

because it determines the amount of compression after the bunch phase space is tilted in the buncher cavity. The length of this second drift is tuned to match from a roughly flat beam at the cathode to a compressed beam with bunch length 13.9 ps, matching the distribution at 40.8 MeV.

In the design by Steve Lidia, the injector acceleration consists of six cavities at 1300 GHz to accelerate the beam and generate an energy chirp. The peak energy gain per cavity starts at 3.07 MeV and rapidly increases to 7.61 MeV. The phase of the first cavity is shifted -5 degrees, the rest are not shifted. The final beam energy is 38.9 MeV ($\gamma \approx 76.15$). In order to define an initial distribution which has as low an energy and energy spread as possible, the first accelerating cavity is altered to have a peak energy gain of 5.01 MeV. The initial distribution then has a bunch length of 62.7 ps, and an average of 0.517 MeV total energy or 6.3 keV kinetic energy. The initial energy profile is mostly parabolic with a shallow curvature, except near the head where the distribution folds over on itself, the kinetic energy goes up to 19.9 keV, and there is also a large jump in energy spread and current. The minimum energy is 3.1 keV, and the rms energy spread is 2.8 keV, although the slice energy spread is much smaller at around 0.21 keV. The current in the core of the beam is approximately 9 A, and peaks in the head of the bunch with a value of 40 A. The spike in the head of the beam is smeared out because the highest energy particles are displaced towards the tail of the bunch. This is about as cold a starting beam as can be obtained given the non-physical modeling for the dynamics at low energy, while still matching into the target beam profile at 40.8 MeV. Note that the nominal initial bunch length used in the reference injector study is 100 ps, but the value used here of 62.7 ps is of the same order and is foreshortened by the fact the a portion of the beam is folded over on itself. The local compression factor in the low-energy core of the bunch is 7.5 going from the initial distribution to the end of the injector, and the chicane compresses the bunch by an additional factor of roughly 12. The initial core current used in the injector models is in the range 6 A to 8 A, fairly close to the value seen here. The spurious starting distribution for the head of the beam is not expected to be a significant problem for jitter studies, especially as the properties of the core of the beam are the main interest for operation in an FEL.

The rest of the beamline has already been tuned to yield a desirable, reasonably flat profile at the final energy of 2.423 GeV with a current of 0.836 kA, except near the endpoints of the beam. This design uses two linac sections, where linac 1 has 18, 1300 MHz cavities set to a phase of -16.6 degrees with a peak energy gain per cavity of 12.97 MeV. This is followed by 7 third-harmonic cavities at 180 degrees with a peak energy gain per cavity of 3.41 MeV. At this point the beam energy is 240 MeV, and the current in the core of the bunch is about 69 A. A bunch compressor follows, consisting of two chicanes each with R_{56}=-67.5 mm, which reduces the bunch length to 0.8 ps. Linac 2 accelerates the beam to the final energy with 162, 1300 MHz cavities each of which has a peak energy gain of 13.5 MeV and a phase of -2.0 degrees. The beamline ends with the "spreader" section which has no net linear R_{56} but introduces a small nonlinear dispersion.

IV. BEHAVIOR AT NOMINAL PARAMETERS

The evolution of the nominal bunch phase space is shown in Fig. 1. Although the starting distribution is not physical, it does match into the desired intermediate phase space. The final distribution is fairly flat, with a slight dip in energy near the head and current spikes at both the head and tail of the bunch. Two important aspects of the beam are the bunch length and the position of the center of the bunch, which is important for proper synchonization. For the nominal parameters, the bunch is 0.265 mm long (0.885 ps) and the midpoint has an offset of -0.0076 mm (-0.025 ps). Through most of this report, longitudinal position is given in terms of t, where positive t corresponds to the tail of the bunch. Where $z = ct$ is used instead, positive z also corresponds to the tail of the bunch, where the particles pass a given location at later times. This follows the convention for LiTrack.

Other properties of the final bunch are analyzed in a region 0.5 ps long. The nominal center of the region of the bunch taken for diagnostics is $t_{nom} = 0.0$ ps, close to the midpoint but shifted a small amount towards the tail. The flat part of the energy and current profiles are shifted towards the tail of the bunch. On the other hand, the tail of the bunch is more sensitive to jitter in parameters. To pick a specific value, the nominal energy is taken as the nominal energy at t_{nom}, $E_{nom} \simeq 2.4232$ GeV, and energies will be given below as deviations from this value in keV. The head and tail of the beam occur at -0.467 ps and 0.417 ps, respectively. The energy at the nominal location has, by definition, no deviation from the nominal energy. Within the 0.5 ps region surrounding the nominal center, the average energy is 9.1 keV below E_{nom}, and the rms standard deviation in slice energy is 20 keV. The current at the nominal center is 0.847 kA, and the average current in the selected region is 0.836 kA. The rms standard deviation in current is 34 A.

The energy and current profiles are shown in Fig. 2, as are smoothed values within the selected region. The smoothing is obtained through an 8th-order polynomial fit. All derivatives quoted below will be taken from the smoothed profiles. As can be seen, the smoothing is particularly important for the current profile. Using smoothed variables, the rms current deviation drops to 25 A, while the rms energy deviation is essentially unchanged. Each
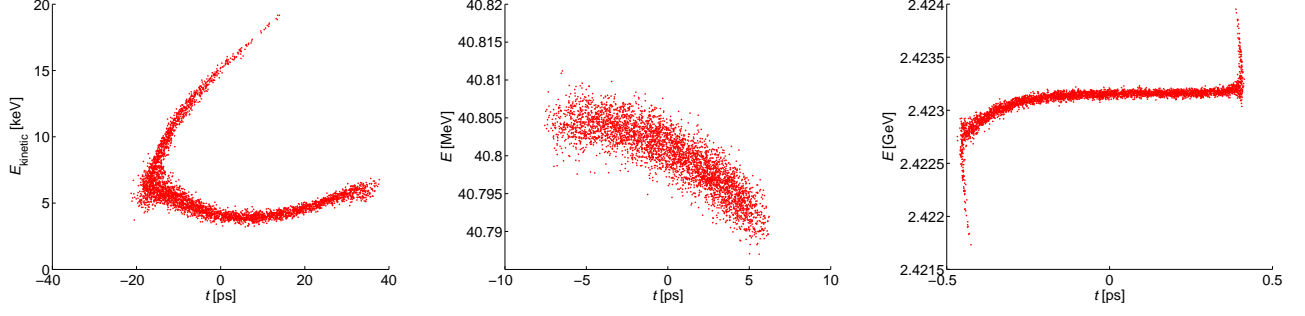
FIG. 1: Phase space distributions: fictional distribution at the cathode (left) matches into the distribution at the start of linac 1 (center), leading to the final beam distribution (right).

bin has roughly 1000 particles, for simulations that have a total of 100,000 macroparticles, so statistical errors of a few percent are to be expected in the current diagnostic. This makes smoothing an important aspect of the analysis, especially for the beam current. The average energy profile exhibits much smaller fluctuations because the energy spread itself is a very small fraction of the beam energy. The energy spread used in these simulations is taken especially low in order to highlight short-wavelength wakefield effects, which tend to be obscured by large energy spreads. For these parameters, the slice energy spread at the nominal center is about 23 keV, which is also the typical value within the 0.5 ps region of interest.
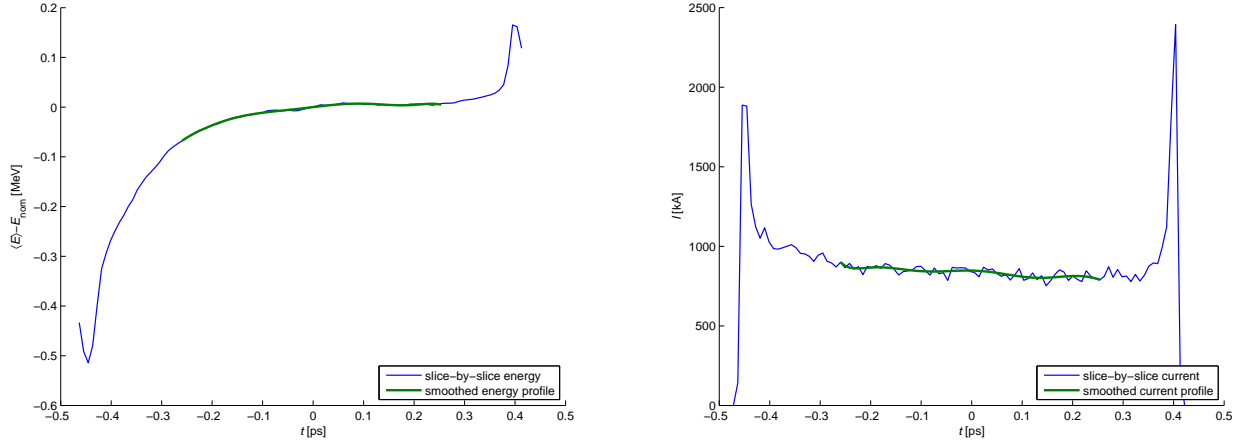

FIG. 2: Energy (left) and current (right) profiles for nominal lattice and beam. Also shown are smoothed profiles in the target interval, -0.25 ps < t < 0.25 ps.

We are also interested in the slope and curvature of the energy and current profiles within the region of interest. At the nominal center, the slope in energy is 0.12 MeV/ps, and the second derivative is 0.0045 MeV/ps^2. Across the entire region of interest, the average energy is lower by 9.1 keV, the rms spread in slice average energy is 20 keV, the rms spread in slope is 0.19 MeV/ps and the rms spread in curvature is 4.2 MeV/ps^2. The derivatives in energy get large towards the front (most negative z) of the selected beam interval. For current, at the nominal center the current is 0.847 kA, the slope is -0.070 kA/ps and the curvature is -8.18 kA/ps^2. Across the entire region of interest, the average current is 0.836 kA, the rms spread in slice current is 25 A (after smoothing), the rms spread in slope is 0.54 kA/ps and the rms spread in curvature is 35 kA/ps^2. There is significant slice-by-slice variation in the instantaneous current which may not be physical, but even after smoothing there is some ripple in the current profile. Simulation parameters and results are summarized in Table I.

TABLE I: Linac design parameters. All beam energies are in terms of total energy.

Quantity	Value
macroparticles	100,000
number bins	100
Initial bunch length	62.7 ps
Bunch length after buncher	13.9 ps
Final bunch length	0.885 ps
Initial peak current in core	0.009 kA
Peak current after buncher	0.069 kA
Final peak current	0.863 kA
Energy after buncher	0.00115 GeV
Energy at chicane	0.240 GeV
Final energy	2.423 GeV
At center:	
energy offset	0 keV
energy slope	0.12 MeV/ps
energy curvature	0.0045 MeV/ps^2
current	0.847 kA
current slope	-0.070 kA/ps
current curvature	-8.18 kA/ps^2
Over 0.5 ps interval:	
average energy offset	-9.1 keV
energy variation	20 keV
energy slope variation	0.19 MeV/ps
energy curvature variation	4.2 MeV/ps^2
average current	0.836 kA
current variation (smoothed)	25 A
current slope variation	0.54 kA/ps
current curvature variation	35 kA/ps^2

V. RESPONSE TO SHIFTS

The sensitivity of this configuration to fluctuations in the RF cavities and chicanes has been examined, first by considering errors in single elements. RF cavities can have errors in power and phase, and the chicane can have an error in R_{56}. Also, the initial bunch generated at the photocathode can have an error in timing and in total charge per bunch. An additional effect, errors in the initial bunch average energy, has been included but barring large drifts in the voltage at the cathode this is not expected to be a realistic type of jitter. The fact that these simulations begin roughly at 6.3 keV kinetic energy instead of closer to rest could conceivably obscure some dynamics that would lead to a jitter in energy at the corresponding distance from the cathode. Timing errors for the creation of the bunch will lead to a phase offset not only for the photocathode cavity, but with all subsequent RF because time of flight should be nearly fixed, especially once the beam becomes highly relativistic. This error will have a strongly correlated effect through the entire beamline. The velocity compression at low energy reduces timing errors by shifting particles towards a fixed RF phase, although it also introduces sensitivity to the voltage and phase of the buncher cavity, and to a lesser extent the cathode RF cavity. The sections used to model drifts are kept fixed. Variations in charge per bunch will also have a coherent effect via the wakefields as the beam passes through all of the RF cavities. Note that any impact of varying charge per bunch related to space charge forces cannot be captured by these simulations in LiTrack.

Because it is complicated to ascribe increases in the spread of energy derivatives to single parameters, we instead focus on the value of the these derivatives at the nominal center. In Table II, changes in central current, energy, and local derivatives of the energy profile at the nominal center are given for a single parameter being shifted away from its nominal value.

VI. GLOBAL JITTER STUDIES: METHODS

We now consider the combined effect of jitter in all quantities except for the initial beam energy, which is not expected to have any physical basis assuming the voltage at the cathode can be held constant. In the nominal case, all RF cavities are taken to have jitter in peak energy gain of 0.01% and jitter in phase of 0.1 degrees, except for the 100 MHz RF at the cathode where the jitter tolerance is relaxed to 0.02% and 0.2 degrees. The values of R_{56} in the

TABLE II: Response of key beam parameters to input errors. All parameters are calculated at the nominal beam center, and measured in terms of smoothed profiles.

Area, Item	Shift	$\Delta I/I_{\mathrm{nom}}$	$\Delta E/E_{\mathrm{nom}}$	$\Delta E'$ [MeV/ps]	$\Delta E''$ [MeV/ps^2]	Bunch Length [fs]	Midpoint [fs]
Cathode							
N_e	1%	-0.0047	-9.8×10^{-6}	-0.054	0.17	13.0	5.89
Timing	100 fs	-6.2×10^{-4}	2.8×10^{-6}	0.0060	0.051	1.81	-0.79
Init energy	100 eV	0.015	-1.9×10^{-5}	-0.0061	-0.30	-11.8	1.50
RF voltage	0.01%	0.0098	-1.4×10^{-5}	-0.0038	-0.19	-8.56	1.07
RF phase	0.1 degrees	0.0069	-1.0×10^{-6}	0.012	0.17	-0.77	-0.35
Buncher							
RF voltage	0.01%	-0.0012	2.1×10^{-6}	-0.0086	-0.14	1.38	-0.13
RF phase	0.1 degrees	0.063	-9.1×10^{-5}	-0.034	-0.34	-56.7	6.51
Injector Accel (6 cavities)							
first RF voltage	0.01%	-2.3×10^{-4}	1.6×10^{-8}	1.1×10^{-4}	-0.025	0.060	-0.93
first RF phase	0.1 degrees	-0.0014	3.4×10^{-7}	-0.0067	-0.29	1.73	-1.50
second RF voltage	0.01%	-2.4×10^{-4}	2.2×10^{-8}	4.2×10^{-4}	-0.035	0.073	-1.01
second RF phase	0.1 degrees	-0.0020	2.5×10^{-7}	-0.0043	-0.032	1.76	-0.083
third RF voltage	0.01%	-2.3×10^{-4}	2.7×10^{-8}	2.2×10^{-4}	-0.036	0.090	-1.22
third RF phase	0.1 degrees	-0.0023	3.1×10^{-7}	-0.0040	-0.040	2.12	-0.10
remaining RF voltage	0.01%	6.5×10^{-4}	1.1×10^{-7}	-0.0025	-0.29	0.103	-1.43
remaining RF phase	0.1 degrees	-0.0025	3.7×10^{-7}	-0.0038	-0.055	2.48	-0.12
Linac 1 (18 cavities)							
RF voltage	0.01%	4.9×10^{-4}	1.2×10^{-7}	-0.0024	-0.31	0.103	-2.33
RF phase	0.1 degrees	-0.0035	1.0×10^{-6}	-0.0066	-0.79	4.94	-12.2
Third Harmonic Cavity (7 cavities)							
RF voltage	0.01%	7.8×10^{-5}	-1.8×10^{-8}	-5.3×10^{-4}	0.014	-0.047	0.63
RF phase	0.1 degrees	0.0033	-3.6×10^{-7}	-0.0032	-0.15	-3.3	0.18
Chicane							
R_{56}	0.01%	0.0015	4.4×10^{-6}	6.3×10^{-4}	0.778	-1.12	22.6
Linac 2 (162 cavities)							
RF voltage	0.01%	-1.4×10^{-7}	5.6×10^{-7}	-3.6×10^{-7}	4.9×10^{-6}	0	0
RF phase	0.1 degrees	-1.1×10^{-7}	3.3×10^{-7}	-1.9×10^{-4}	1.6×10^{-5}	0	0
Spreader							
T_{566}	0.01%	1.5×10^{-4}	-1.4×10^{-10}	-2.9×10^{-6}	-3.6×10^{-5}	-0.027	0.0067

chicane and T_{566} in the spreader are assumed to fluctuate by 0.01% as well. The initial timing error is set to 300 fs, and the rms fluctuations in charge per bunch is 2%. We also consider a more conservative set of errors where the charge per bunch fluctuates by 5%, and most RF phase errors are allowed to be as large as 0.2 degrees except for the buncher RF which is kept at 0.1 degrees. The peak energy gain and phase errors across a single RF cavity are taken to be uncorrelated with each other. The fluctuations of separate elements are also assumed not to have any correlations in this and the following section. As noted above, however, charge per bunch and initial timing errors will propagate throughout the beamline in an especially consistent way.

Rather than rewriting the input file for each instance of a jittered run, a call to a script file is inserted in the LiTrack program after the beamline description file is read in, which modifies the input parameters according to the input seed number. The input seed number is used to label the sample number of a set of runs with jitter rather than directly as a seed to a random number generator. The rms variation of each parameter is defined in this script file, based on the code number of the element and, for RF sections, the wake field file associated with it. Chicanes and generalized dispersive sections have a single error parameter which scales the dispersion and higher order terms by the same factor, and RF sections have independent errors in maximum voltage and phase. Standard chicanes already utilize a special syntax for R_{56} errors which includes the temporal shift in the beam due to a change in the delay through the chicane, in other words the fact that electrons at a different energy will be delayed by the expected amount. More general dispersive sections can be defined by the values of R_{56} and nonlinear terms at a given nominal energy; the longitudinal position of a particle at this nominal energy is taken as unchanged after passing through the dispersive section. In this more general case, we use a simple prescription to account properly for shifts in longitudinal position, that in order to increase the chicane R_{56} by $1 + \alpha$ at the nominal energy, it is necessary to simultaneously increase both the nominal R_{56} and the nominal energy by $1 + \alpha/2$. The nonlinear dispersion terms are multiplied by this same factor. An additional change is made to the LiTrack function in order to extract a summary of information about

the final beam for subsequent diagnostics. For this purpose, a call to an additional script file is inserted at the end of the function. Information such as the sample number, nominal energy, and upper and lower bounds on position and energy are written to an output file, as are histograms of slice current, average energy offset and energy spread, with bins in longitudinal position. For the spreader section, even though T_{566} is allowed to vary, we do not consider fluctuations which lead to having a nonzero net R_{56}.

Given the very large number of independent variables that have been allowed to fluctuate (395 including the initial energy error, which is implemented in the simulations but set to have 0 rms), a purely random distribution of errors may yield reasonably accurate rms values but will result in noise in the statistical analysis unless an extremely large number of sample runs is simulated. To avoid this, the errors in each simulation are not completely random, and restrictions have been imposed on the errors over the entire sample set. First, for each source of error, the mean error is forced to be zero and the rms is forced to have the nominal value. Also, the set of errors from each source over all sample runs, viewed as a vector, is required to be orthogonal to the vector of errors from every other source. This is performed by a simple orthogonalization procedure, where each error source in turn has its vector of errors projected onto the space orthogonal to the error vectors from all previous sources, and is then rescaled to its nominal rms value. Because each error vector has to average to zero, there is one lost degree of freedom and the number of samples must be at least one greater than the number of error sources. In these examples, we take the minimum number of one greater, or 396 samples. We initially set an upper bound of 2.5 standard deviations for each error, but the renormalization process increases this to 2.8 times the actual rms error. The diagonalization process imposes additional constraints that occasionally force us to consider larger errors for some sample runs, of up to 4.2 standard deviations, if we are to keep the nominal standard deviation for each error term. This problem becomes noticeable in the last elements of the beamline where the orthogonality condition requires more adjustment to the errors. In fact, excluding linac 2 the maximum error modelled is 3.2 standard deviations. Running more simulations would reduce this effect, but we are not concerned about occasional large errors in linac 2 because of the very large number of elements in linac 2, each of which have an identical effect on the beam.

This methodology is a simplified alternative to orthogonal array testing [7]. The simplest version of orthogonal arrays, having strength 1, is similar to a Latin hypercube array [8], meaning that for each error, the allowed range is divided into regions of equal probability, and each region is equally represented in the numerical data set. This guarantees uniform sampling of the possible errors, and reduces variability from having a small sample size. However, it is still possible for pairs of elements to have errors which appear correlated due to the small sample size. Orthogonal arrays with strength 2 eliminate all such pairwise correlations while preserving the uniform sampling of each error. Here, because of the large number of elements and sample runs, we are more concerned with eliminating spurious correlations between elements than achieving exact uniformity over each individual element. Thus, we simplify the orthogonal array scheme to eliminate the Latin hypercube feature, and instead emphasize the elimination of all correlations and require the error vectors from different elements to be exactly orthogonal. Also, each individual error vector averages to zero and has exactly the nominal standard deviation.

For any individual sample run, the errors from different sources will appear completely uncorrelated and unconstrained. In particular, there is no restriction preventing unbalanced errors across all cavities, they will merely be statistically unlikely. There is some effect on the exact statistical behavior of the last error sources, as has been already mentioned, but the mean and rms are forced to be at the nominal values. Over all sample runs, the set of errors from any one source will define a unique basis vector against which any output criterion can be projected, yielding an estimate of how much of an impact each source of error had on that output. While this is an accurate and convenient method for determining correlation parameters between error sources and output figures of merit, the main danger of this method is that any nonlinear behavior will lead to spurious correlations as nonlinear effects from one source will be mapped into other degrees of freedom that are associated with other sources. The only real cure for this is to have enough samples that most degrees of freedom are not associated with any error source, and spurious correlations can then be noted by the lack of any corresponding source term. However, this would again require a much greater number of simulations, although possibly not as many as a purely random set of errors would require. In any event, in this work we will use the minimum of 396 samples and simply analyze the results as if all nonlinearities are weak. A weak check on these results can be taken from the initial energy error, which has been suppressed in the calculations, so any correlation of the results with that corresponding error vector indicates a failure in the method, although one single additional degree of freedom is not sufficient to verify the results. An additional consistency check is to compare the response of the beam to errors in different cavities in linac 2; in theory the response function should be identical as the longitudinal position of electrons is frozen throughout linac 2 because of the high beam energy. Note that by construction, the sum of the squares of the correlation parameters between any given output and all of the sources of error will be exactly unity.

VII.   GLOBAL JITTER STUDIES: RESULTS

The nominal jitter in parameters for this analysis consists of 2% in charge per bunch, 300 fs in photocathode timing, 0.01% in chicane strength, 0.01% in RF power and in 0.1 degrees in RF phase, except for the cathode RF which is allowed to have double the power and phase error. Profiles of runs for the case of nominal jitter are shown in Fig. 3. Results for when the jitter constraints are relaxed to 5% in charge per bunch and 0.2 degrees in RF phase (except for the RF buncher which is held at 0.1 degrees jitter in RF phase noise) are shown in Fig. 4. The current profile is very stable near $t = 0$, except when the head or tail of the beam (usually the tail) encroaches on this region. The energy profile mostly changes in two ways from its nominal flat core with a dip in energy at the head and an increase at the tail: either the slice energy peaks close to the head and then decreases linearly towards the tail, and at the same time the beam is shifted towards the tail; or the increase in slice energy near the head continues and only gradually flattens out towards the tail, and at the same time the beam is shifted towards the head. In the first case, the beam energy at the tail is significantly lower than expected, and in the second case the beam energy at the tail is significantly higher.
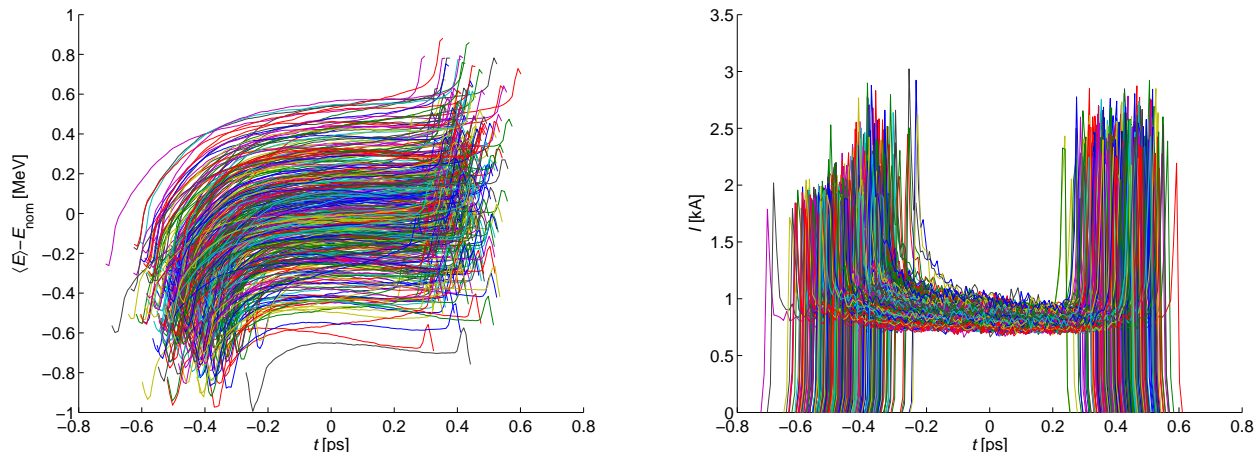
FIG. 3: Energy (left) and current (right) profiles for simulations with different examples of errors, such that the rms charge error is 2%, the rms timing error is 300 fs, and the phase errors are 0.1 degrees except for the cathode RF which is set to 0.2 degrees. This corresponds to nominal jitter levels.
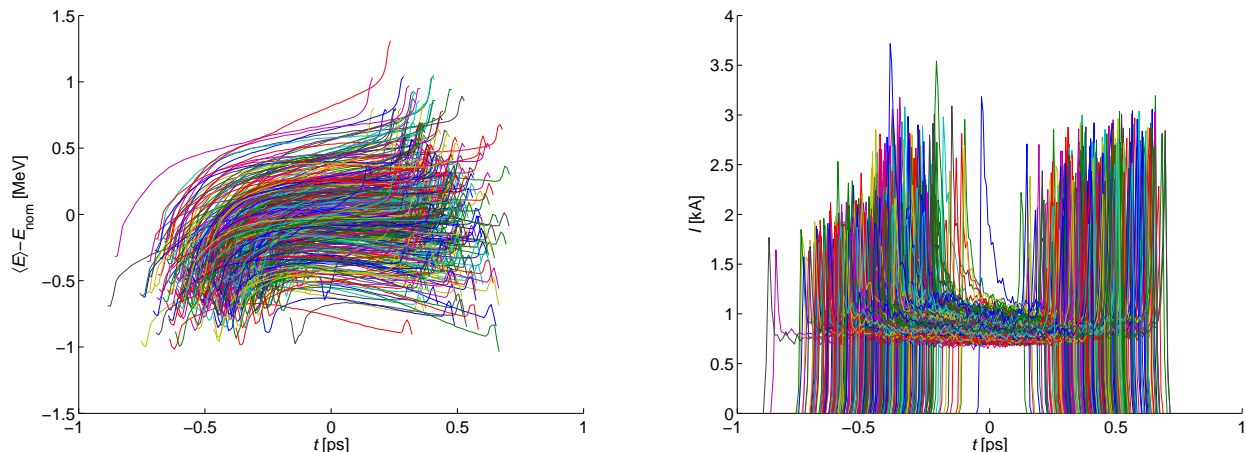

FIG. 4: Energy (left) and current (right) profiles for simulations with different examples of errors, corresponding to relaxed jitter constraints. The rms charge error is 5%, the rms timing error is 300 fs, and the phase errors are 0.2 degrees except for the RF buncher which is kept at 0.1 degrees.

The results of these global jitter studies are summarized in Table III. Some results are given in terms of local properties at the nominal center of the beam, others are averaged over a 0.5 ps region. The average values are

calculated by ignoring any portions of the -0.25 ps < t < 0.25 ps region which lie outside of the bunch itself. In addition, jitter in the bunch length and arrival time are indicated. Four cases are considered, including the nominal jitter case and a 0.5 ps region for averaging. The conservative example is also studied with a charge per bunch jitter of 5% and linac phase errors of 0.2 degrees except for the RF buncher, which is kept at 0.1 degrees. The nominal jitter case with no bunch charge fluctuations at all is also looked at; because the bunch charge only acts on the beam through wakefields, this is reflective of the jitter for a low-charge configuration where wakefields have a much weaker effect. Finally, the nominal jitter case is re-evaluated for a narrower time window of 0.3 ps, with -0.15 ps < t < 0.15 ps, instead of 0.5 ps as for the other 3 cases. The same smoothed profiles as for the original nominal case are used when performing calculations for the 0.3 ps time window.

TABLE III: Fluctuations in output beam parameters under different input jitter conditions and with different time windows.

		Input jitter levels (buncher RF fixed at 0.1 degrees):			
		5% charge	2% charge	2% charge	0% charge
		0.2 degrees RF	0.1 degrees RF	0.1 degrees RF	0.1 degrees RF
	Unit	0.5 ps window	0.5 ps window	0.3 ps window	0.5 ps window
Fluctuations at nominal center:					
energy	MeV	0.27	0.24	0.24	0.23
energy slope	MeV/ps	0.42	0.13	0.13	0.064
energy curvature	MeV/ps^2	3.3	1.4	1.4	1.4
current	kA	0.10	0.058	0.058	0.057
current slope	kA/ps	2.1	0.30	0.30	0.30
current curvature	kA/ps^2	57	9.9	9.9	9.3
Values across target region:					
jitter in avg energy	MeV	0.27	0.24	0.24	0.23
typical spread in energy	MeV	0.052	0.028	0.013	0.025
rms of spread in energy	MeV	0.036	0.016	0.0080	0.014
min spread in energy	MeV	0.0050	0.0048	0.0029	0.0063
max spread in energy	MeV	0.22	0.11	0.041	0.10
typical spread in energy slope	MeV/ps	0.55	0.25	0.099	0.0.25
rms of spread in energy slope	MeV/ps	0.71	0.24	0.056	0.24
min spread in energy slope	MeV/ps	0.058	0.078	0.033	0.0.079
max spread in energy slope	MeV/ps	4.0	2.4	0.59	2.5
typical spread in energy curvature	MeV/ps^2	22	5.4	1.4	5.6
rms of spread in energy curvature	MeV/ps^2	48	12	1.5	13
min spread in energy curvature	MeV/ps^2	1.4	1.2	0.61	1.2
max spread in energy curvature	MeV/ps^2	3200	170	21	170
jitter in avg current	kA	0.090	0.062	0.058	0.063
typical spread in current	kA	0.080	0.043	0.033	0.044
rms of spread in current	kA	0.086	0.033	0.0067	0.036
min spread in current	kA	0.019	0.020	0.020	0.021
max spread in current	kA	0.53	0.33	0.075	0.33
typical spread in current slope	kA/ps	3.2	0.81	0.25	0.83
rms of spread in current slope	kA/ps	6.4	1.7	0.23	1.8
min spread in current slope	kA/ps	0.15	0.17	0.084	0.14
max spread in current slope	kA/ps	58	18	2.5	18
typical spread in current curvature	kA/ps^2	210	48	7.5	49
rms of spread in current curvature	kA/ps^2	510	100	11	110
min spread in current curvature	kA/ps^2	4.2	4.3	1.9	4.5
max spread in current curvature	kA/ps^2	5400	1200	120	1200
Fluctuations in global bunch properties:					
bunch length	fs	99	69	69	64
arrival time	fs	110	59	59	58

The spread or jitter in derivatives of the energy and current profiles appear to be quite large, of order 1 MeV/ps and 1 kA/ps, respectively. A few of the most extreme cases arise when the beam tail overlaps the nominal region of interest, in which case very high derivatives are expected and the FEL output can be drastically altered. In most cases, however, the jitter in slope corresponds to less severe distortions. The effect on the FEL output can be roughly estimated as a relative frequency error of $2\pi R_{\text{FEL}}(1/cE)dE/dt$, where c is the speed of light, R_{FEL} is the R_{56} of elements in the FEL section and dE/dt is the energy slope. Taking a large value of 100 microns for R_{FEL} yields a relative frequency error of 2.5×10^{-4}. This corresponds to 300 meV for a 1 nm wavelength, which as a spread in

frequencies corresponds to a coherence length of 25 microns (0.08 ps). The spread in slopes is a major criterion for how close the FEL output can be to the bandwidth limit. The corresponding values for energy curvature is larger, with a worst-case scenario of a factor of 10 larger spread in frequencies for a 100 micron duration pulse, but a more likely consequence is that these large spreads correspond to shorter-period distortions and will yield other types of undesirable FEL output such as distinct sidebands or a complex structure of the power profile versus time. For current, a typical variation in current of 4% over the region of interest will lead to power variations and some phase error as well. For the current slope, a spread of 1 kA/ps corresponds to a 30% variation over 0.25 ps. Typically, this will result in a phase shift of approximately 1 radian or less, although this will depend on the specific FEL design. Therefore, jitter in the current profile is expected to be most apparent in the power profile, although a short-wavelength oscillation in current could lead to the generation of frequency sidebands.

The fluctuations in average energy, energy slope, and current all seem to be dominated by the phase error of the RF buncher. RF phase errors in linac 1 are slightly more important for jitter in the slope in current, are also significant factors in the average current and energy slope, and are the main contributor to jitter in the arrival time. The bunch length is impacted by the charge per bunch, the buncher phase, and the phases in linac 1. Reducing the length of the targeted portion of the beam from 0.5 ps to 0.3 ps significantly improves the beam uniformity although the jitter in average energy is not improved at all.

The phase of the buncher cavity has a large impact on the final beam. This is because the beam is compressed towards the zero phase of the RF, and a 0.1 degrees error corresponds to a shift in the zero phase of 430 fs. Given the compression factor of roughly 7.5, the displacement of the compressed bunch is only slightly smaller at 370 fs. On the other hand, the error in timing for production of the bunch at the cathode is itself compressed by the RF buncher, so that a 300 fs initial timing error corresponds to a timing error in the compressed bunch of only 40 fs. Differences in the energy gained by the bunch as it passes through the cathode also lead to a significant shift in time by the end of the drift section after the RF buncher. Given the 0.85 m of drift between the buncher and the accelerating cavities, an error in kinetic energy of 0.02% leads to a timing error of roughly 100 fs after compression. The second bunch compression using a magnetic chicane acts similarly, compressing the bunch by a further factor of 12.

TABLE IV: Relative contributions to jitter from the most important sources, all in %. The quantities considered do not make any reference to the time interval used for calculating averages.

	Input jitter levels:		
	5% charge	2% charge	0% charge
	0.2 degrees RF	0.1 degrees RF	0.1 degrees RF
Central energy			
charge	22	4	-
cathode voltage	6	8	8
bunch phase	68	86	90
Central current			
charge	2	2	-
cathode voltage	3	6	6
buncher phase	28	68	70
linac 1 phase	39	17	18
third harmonic phase	2	2	2
Bunch length			
charge	42	14	-
cathode voltage	3	6	7
buncher phase	32	67	78
linac 1 phase	18	9	11
third harmonic phase	3	2	2
Arrival time			
charge	7	4	-
linac 1 voltage	< 1	3	3
linac 1 phase	87	77	80
chicane	4	14	15

The relative importance of different sources of jitter for the central energy, central current, bunch length and bunch midpoint is shown in Table IV. The buncher RF phase is the dominant contributor to jitter in the central energy; when the charge per bunch variations are increased to 5%, the charge becomes a significant source of energy jitter but still several times smaller than the buncher RF phase. The voltage of the cathode RF also contributes somewhat to energy jitter. The nature of the fluctuations induced by the buncher RF errors is that the bunch profile is not greatly altered, but the entire bunch is either shifted up in energy and made longer, or shifted down in energy and

compressed. In addition, when the overall energy is lower, the energy in the core of the beam tends to be chirped and decreases towards the tail of the beam. The location of the head of the beam jitters much more than the location of the tail. The RF phase of the buncher is also the main contributor to jitter in the central current except when the linac RF phases are allowed to fluctuate by 0.2 degrees, in which case the linac 1 phases become the dominant term. Jitter in the bunch length is dominated by charge errors when these are at the 5% level, but dominated by the phase of the buncher RF otherwise. Higher charge per bunch leads to longer bunch lengths, to the extent that the expected tendency to have more peak current when there is more charge per bunch is in fact reversed. The arrival time jitter is dominated by the phases in linac 1, with a smaller contribution from fluctuations in the strength of the chicane. The effect of charge per bunch fluctuations also becomes noticeable when these are at the 5% level. Large jitter in the timing of the core region of the bunch may reduce the "useful" part of the bunch for practical FEL operation to below the 0.5 ps width taken for most of the examples. Charge per bunch, because of the slope of the wakefield forces, also has the largest effect on the slope of the energy profile, as can be seen from Table II.

TABLE V: The contributions of various sources of jitter to fluctuations in output average beam energy, average current, bunch length, and the shift in arrival time. Energy and current are averaged over a fixed 0.5 ps interval. The sign represents the sign of the perturbation when the input error is positive.

Quantity	Amount	Contribution to Jitter in:			
		Energy [keV]	Current [A]	Bunch length [fs]	Arrival time [fs]
charge	1%	-25.5	-1.9	11.3	6.0
timing	100 fs	8.5	-0.68	1.9	-0.83
cathode voltage	0.01%	-33.3	8.1	-8.7	1.1
cathode phase	0.1 degrees	-2.1	2.2	-0.73	-0.43
buncher voltage	0.01%	4.2	-1.5	1.3	-0.15
buncher phase	0.1 degrees	-220	50.7	-56.3	5.7
injector accel voltage	0.01%	0.83	1.3	0.14	-3.0
injector accel phase	0.1 degrees	1.8	-5.2	5.3	-0.93
linac 1 voltage	0.01%	0.76	-1.1	0.43	-10.0
linac 1 phase	0.1 degrees	11.5	-26.7	21.0	-51.7
third harmonic voltage	0.01%	-1.3	1.1	-0.14	1.7
third harmonic phase	0.1 degrees	-3.4	7.1	-8.7	0.53
chicane	0.01%	9.7	3.8	-1.1	22.3
linac 2 voltage	0.01%	16.8	0.092	0.0033	-0.037
linac 2 phase	0.1 degrees	10.1	1.5	-0.063	0.13
spreader	0.01%	-0.43	1.4	-0.016	0.13
null	-	0.13	0.13	0.024	0.033

The amount of jitter in average energy and average current contributed by various sources are delineated in Table V. Note that the total jitter can be determined by scaling these contributions to the actual input fluctuations, and adding independent terms in quadrature. The "null" example represents a fictitious error source that does not actually affect the particle dynamics, and is a rough estimate of the accuracy. The values given are calculated from the case where the charge per bunch, timing, and RF phases have the nominal variability, but are reasonably consistent with the other cases. The calculations for the linac 2 voltage and phase have to be done carefully because of the large number of independent cavities each of which have a statistical uncertainty which is almost certainly larger than the true effect. Instead of accepting at face value the correlations from all of the cavities, we use the fact that each cavity in linac 2 should have the same response function in terms of the effect of a given voltage and phase error on the beam, as seen in Section V. In contrast, the estimated correlations for neighboring cavities can have swings in value that are large enough that they can even change sign, especially when considering correlations with the average current. Therefore, we assume that all cavities in linac 2 must have the same correlation, and average the correlation parameter to obtain a better estimate of the correct value. The contribution of all 162 cavities in linac 2 to fluctuations is then taken to be $\sqrt{162}$ times the average contribution, assuming that the cavities will add noise independently. This procedure substantially reduces the estimated contribution of voltage errors to jitter in the average current, from 8.4 A to a revised value of 0.092 A, and similarly for phase errors, from 9.0 A to a revised value of 1.5 A. The contribution to energy jitter is also slightly reduced in this way, from 17.6 keV to 16.8 keV in the case of voltage errors, and from 11.4 keV to 10.2 keV in the case of phase errors. This revision is necessary because the effect of an error from a single cavity in linac 2 is much smaller than the statistical uncertainty as indicated by the "null" contribution to jitter, which by definition ought to be 0, while the large number of cavities amplifies the effect of this uncertainty. Collective contributions from linac 1, which has 18 cavities, are also slightly corrected by this effect, with the current jitter from voltage fluctuations reduced from -2.9 A to -1.1 A, and from -27.0 A to -26.7 A for phase fluctuations.

Energy jitter from the voltage fluctuations are reduced from 1.7 keV to 0.76 keV, and from 11.7 keV to 11.5 keV for phase fluctuations. The bunch length and arrival time jitter which result from linac 2 errors are also exaggerated, and a similar correction reduces the contribution to arrival time jitter from 0.50 fs to 0.0033 fs for voltage errors and from 0.50 fs to -0.063 fs for phase errors. Similarly, the arrival time jitter contribution is decreased from 0.77 fs to -0.037 fs for voltage errors, and from 0.87 fs to 0.13 fs for phase errors. Other beamline elements give similar results for the two calculation methods. The lack of agreement between different calculations for linacs 1 and 2 gives another indication of the level of uncertainty associated with the statistical methods used here, as does the fact that errors in linac 2 should have negligible effect on the bunch length and arrival time.

As is clear from Tables IV and V, the phase errors in the buncher RF cavity produce outsized amounts of jitter compared to other errors, with the exception of the arrival time which is influenced primarily by the phases in linac 1 and also by the strength of the chicane. The overall energy jitter, which ranges from 240 keV to 270 keV, is almost completely set by the variations in the phase of the buncher RF. Although the jitter in average energy is not reduced by considering a narrower interval for use in the FEL, many other concerns such as the variations in energy slope, current slope, and higher derivatives are greatly improved in this way.

VIII. TIME EVOLUTION AND SPATIAL CORRELATIONS

To further explore the variability of the electron beam parameters as it enters the undulator section, we replace random shot noise with an evolving set of errors that has a characteristic power spectral density (PSD) function. Each error function is generated by a series of Fourier components in 1 Hz bins, ranging from 1 to 300 Hz. Based on experimental observations of acoustic noise [9], the power spectral density is assumed to be flat over the range 1 to 20 Hz, and then decreases at higher frequencies as $1/f^4$, where f is the frequency. The total contribution of frequencies above 300 Hz is sufficiently small that applying a cutoff at this point has negligible effect. Other models for the ground motion include the "ATL" law [10], which sets the power spectral density to scale as $1/f^2$ at low frequencies, leading to a diffusive motion of the displacement, as well as variations including the possibility of shifting to $1/f^4$ scaling at higher frequencies [11]. Normalized to the same rms displacement, the noise model used in this paper with a constant spectrum at low frequencies adds more weight to the power spectral density curve at higher frequencies, and is thus the more pessimistic choice.

In each plot below, the curves represent single shots separated by 1 ms in time. A total of 250 shots, corresponding to 250 ms of beam variation, is shown. This duration is sufficiently short that the low-frequency cutoff at 1 Hz should not introduce any systematic errors, while it is long compared to typical feedback response times and would comprise a large number of total shots at the targeted repetition rate of 100 kHz or above. The PSD is normalized so that the the rms of the combined error from all frequency bins yields the nominal jitter.

We have considered two versions of the correlated noise model. In one, each source of error is completely independent in its variations over time. In the other, most RF cavities are grouped into modules, and in the portion of the linac above 40 MeV, all cavities within a single module are assumed to have identical power errors and identical phase errors (there are never any correlations between power and phase errors). Neighboring modules do not have any correlations between them. The current design has 9 RF cavities in most cryomodules, and if these RF cavities are powered by a single klystron, then strong correlations might be expected. Assuming strong correlations is a more pessimistic case, because there is less averaging out compared to fully independent errors, but simulations do not exhibit significantly worse performance, so in Fig. 5 we only show results for the strongly correlated case.

All sources of jitter are taken to have a similar PSD function except for the charge per bunch. Because there are no external drivers affecting the quantum efficiency of the photocathode, because jitter in the power of the incident laser is expected to be driven by highly nonlinear effects, and because the statistics may be sensitive to specific design choices of the cathode material and laser, the charge per bunch is taken to have purely random errors with no correlations from shot to shot. The effect of variation in the charge thus stands out compared to other sources, especially over short time scales. The charge per bunch variations also have a distinctive signature, as they affect longitudinal wakefields (recall that space charge effects are not modelled by LiTrack). Therefore, the head of the bunch is not affected but the tail of the bunch is very sensitive to charge per bunch.

It is illustrative to examine the first 10 ms of time samples, as we can see that the range of errors is much smaller over this smaller time scale. Although the charge per bunch still introduces random errors, other sources of error change little over this time scale. This suggests that feedback on this time scale could significantly improve the total jitter in the final electron beam.
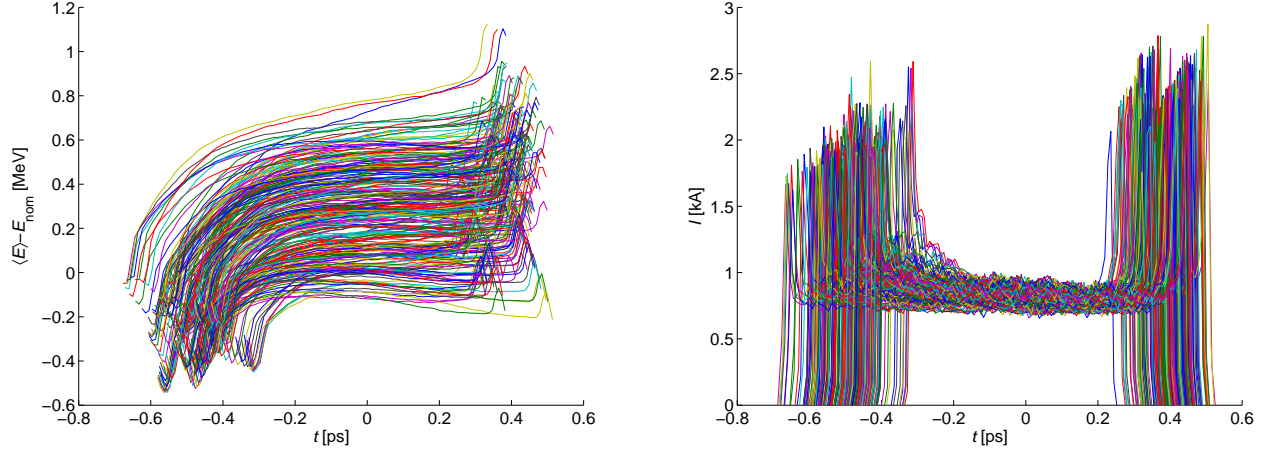
FIG. 5: Energy (left) and current (right) profiles for simulations of sample electron bunches over a 250 ms time period, adopting the noise spectrum described above. The rms charge error is 2%, the rms timing error is 300 fs, and the phase errors are 0.1 degrees except for the cathode RF which is kept at 0.2 degrees. RF cavities within the same cryomodule are assume to have identical errors.
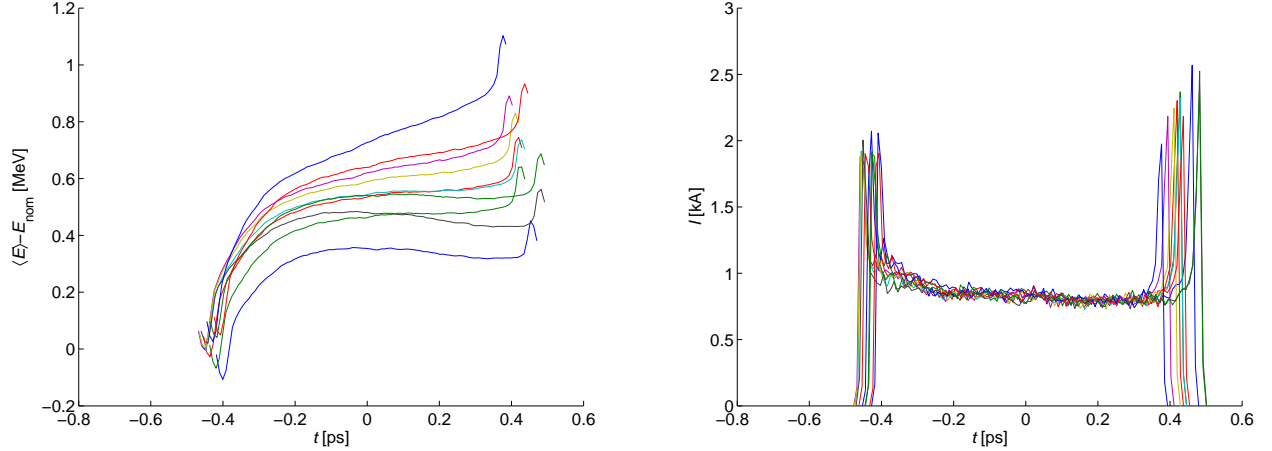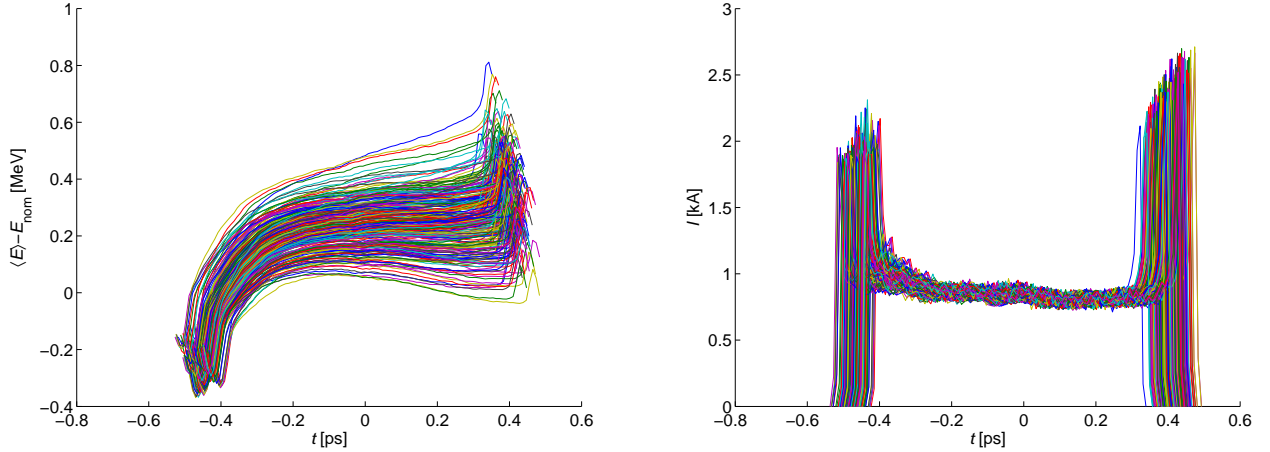

FIG. 6: Energy (left) and current (right) profiles for the example used in Figure 5, but showing only a 10 ms time period.

IX.   TIME EVOLUTION WITH FEEDBACK

Finally, we adopt a simple model for errors which evolve in time and which are corrected by a feedback scheme. The most important beam properties to correct for are electron beam energy, timing, bunch length, and energy chirp. Charge per bunch, as stated above, cannot be corrected for as there are no shot-to-shot correlations. The importance of stable beam energy is to maintain the resonance condition in the FEL, otherwise the gain length and power output could suffer large fluctuations with electron energy. The timing and bunch length jitter should be controlled to maximize the interval of time over which electron beam can be reliably seeded by a laser to produce good output in the FEL. For the ideal configuration, there is approximately 0.2 mm or 650 fs over which the beam current, energy, and energy spread are all close to nominal parameters. Finally, energy chirp in the beam can lead to shifts in the wavelength of the output radiation, although this is more of a concern for long pulses where a narrow bandwidth is desired. All of these quantities could potentially be measured and corrected for by feedback, perhaps through special-purpose RF cavities that have a lower Q factor enabling rapid response times. The feedback could be applied to the final beam or earlier in the acceleration stage, and some errors can be measured at the final bunch compression stage. The energy chirp might be a challenge to measure, and there are higher-order perturbations to the beam which would not be fully corrected by a feedback system. In order to gain some understanding of how a feedback system

could improve performance, we consider a generic feedback system rather than modeling the feedback directly. We apply the general rule that a well-designed feedback system should yield similar results as if the noise inputs were all reduced by some factor. We model the system with feedback by taking the PSD for each noise source and reducing it by the function $1 + (\sqrt{N}/10Tf)$, where f is the frequency, T is the time resolution of the feedback system, and N is the number of measured quantities that the feedback system reacts upon. Here $N = 3$, for example bunch energy, arrival time and bunch length could be measured. The time resolution T of the feedback system is taken to be either 1 ms or 0.1 ms. The feedback does not act on the errors in the bunch charge, which are assumed to vary shot-to-shot.

This feedback model dramatically improves performance, with the charge per bunch variation remaining as the main driver of jitter in electron beam parameters. The nominal rms jitter of 2% in the charge per bunch yields variations in the tail of the bunch which are just above the contribution from remaining sources of error, as shown in Fig. 7. Examples without any charge per bunch fluctuation are shown in Fig. 8. Note that only charge per bunch errors have been reduced here and not the timing error, because the timing error is not very deleterious and its effect is already being reduced by the feedback model.

FIG. 7: Energy (left) and current (right) profiles for simulations of sample electron bunches over a 250 ms time period, assuming a typical noise spectrum and with a generic feedback model with 1 ms time scale. The rms charge error is 2%, the rms timing error is 300 fs, and the phase errors are 0.1 degrees except for the cathode RF which is set at 0.2 degrees. RF cavities within the same cryomodule are assume to have identical errors.

When the jitter in the charge per bunch is at or above 2%, and is varying at a rapid rate, then there is very little improvement to be gained by improving the feedback system. Below, results are shown for the same system as Fig. 7 but with a feedback time scale of 0.1 ms instead of 1 ms. Although the head of the bunch is extremely stable, variations in the tail of the bunch are not much improved and the total time interval over which a laser pulse can reliably be made to overlap the core of the electron bunch is only slightly lengthened. The jitter in energy of the nominal time slice of the beam is reduced from 75 keV to 45 keV. This lower value derives predominantly from fluctuations in the charge per bunch. Thus, there is little motivation to enhance the feedback design further unless charge fluctuations are determined to be less problematic than considered here, or unless the RF errors are significantly worse than adopted here.

Qualitatively, the charge per bunch jitter is the main driver of fluctuations in the energy chirp. For the case without feedback, the timing fluctuations of the head of the bunch are equal to that of the tail, unlike what would be expected from charge per bunch contributions. The jitter with no feedback is sufficiently large that linear perturbation models do not yield accurate results, but with feedback the overall errors are reduced and the linear model is in fair agreement. For the 1 ms feedback model, the bunch length errors are larger than the arrival time errors, in agreement with the first row of Table V. For the 0.1 ms feedback model, the charge fluctuations are the dominant effect. The timing variations of the head of the bunch are reduced to 8 fs, and the tail and bunch length jitter are 25 fs; the arrival time error lies in between these values at 15 fs. The fluctuations in beam energy, current, bunch length, and arrival time are given in Table VI, along with estimates of the contribution of charge fluctuations. Repeating the simulations without any charge per bunch jitter, fluctuations in the energy could be reduced to as low as 23 keV, in the current to 7 A, in the bunch length to 7 fs, and in arrival time to 5 fs.
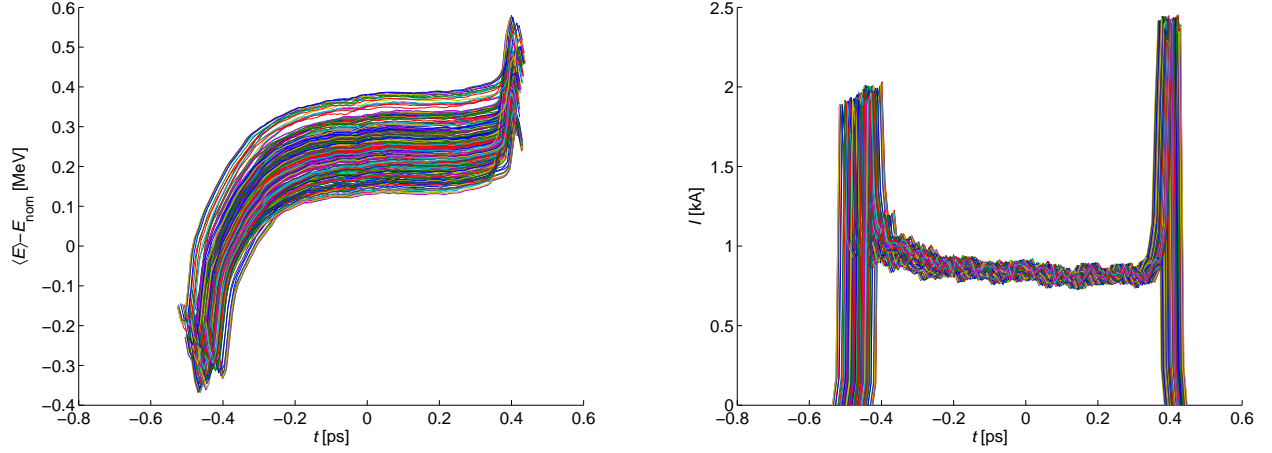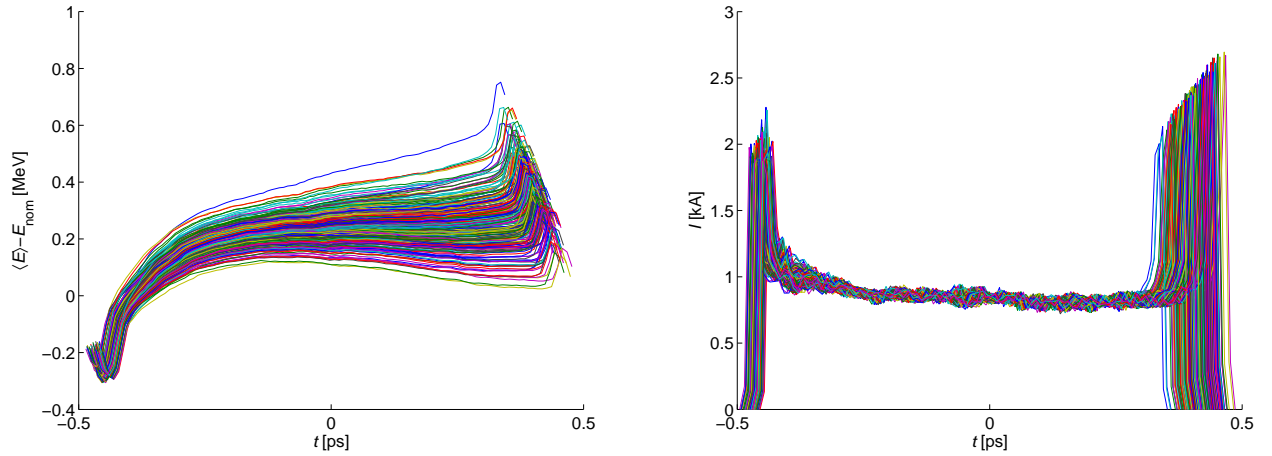
FIG. 8: Energy (left) and current (right) profiles for simulations of sample electron bunches over a 250 ms time period, assuming a typical noise spectrum and with a generic feedback model with 1 ms time scale. There is no rms charge error, which would in practice be the effect of going to a reduced bunch charge where wakefields have a much smaller effect, the rms timing error is 300 fs, and the phase errors are 0.1 degrees except for the cathode RF which is set at 0.2 degrees. RF cavities within the same cryomodule are assume to have identical errors.

FIG. 9: Energy (left) and current (right) profiles for simulations of sample electron bunches over a 250 ms time period, assuming a typical noise spectrum and with a generic feedback model with 100 microsecond time scale. The rms charge error is 2%, the rms timing error is 300 fs, and the phase errors are 0.1 degrees except for the cathode RF which is kept at 0.2 degrees. RF cavities within the same cryomodule are assume to have identical errors.

X. CONCLUSIONS

Jitter in the linac can have a deleterious effect on FEL performance, predominantly through two mechanisms. The time interval over which a laser pulse can reliably be made to interact with the core of the electron bunch can be reduced far below the nominal interval of 650 fs. In addition, the average energy in this core region can fluctuate by an amount comparable to the interaction bandwidth, which will significantly alter the output power for a seeded FEL. Additional fluctuations can come from variations in the current or from an energy chirp across the electron beam.

The generic analysis adapted in most of this work is useful in its own right to evaluate the beamline as a whole, identifying elements to which the final electron beam is especially sensitive and characterizing the form of the resulting beam variation. This model has been further specialized by considering expected magnitudes and temporal structures in the errors in beamline elements. These assumptions should be refined according to experimental data as it becomes available. The feedback model offers quantitative results concerning an important benefit that CW operation allows.

TABLE VI: Jitter in beam energy and current at the fixed reference time, and in the bunch length and arrival time. Results are given for the nominal rms errors and PSD function, for the cases with no feedback, feedback with a 1 ms resolution, and with a 0.1 ms resolution. Includes estimates for the contribution of bunch charge fluctuations, where possible.

Parameter	no feedback	1 ms feedback	0.1 ms feedback
Energy [keV]	200	80	47
contrib from N_e	58	45	41
Current [A]	47	20	11
contrib from N_e	-	9	9
Bunch length [fs]	61	31	25
contrib from N_e	5	23	24
Arrival time [fs]	54	22	15
contrib from N_e	17	15	15

Having a diagnostic for delivered beam errors could allow for corrections through use of the results of Table II to identify changes in the linac which would correct for key errors. In particular, the energy jitter, which is expected to be a dominant feature of the electron beam delivered by the linac, can be countered by dedicated RF cavities. Adjustments in the chicane can also help to correct for timing or bunch length jitter. From the examination of a generic feedback system, it seems that most sources of error can be controlled to acceptable levels with the possible exception of fluctuations in the charge per bunch. The charge per bunch error should be characterized more carefully to determine if they are sufficiently correlated to apply corrections through feedback systems. Hardware studies should also be undertaken to assess the size of the errors which are likely to be seen in practice.

Taking nominal values of 2% jitter in the charge per bunch, RF power and chicane fluctuations of the order of 10^{-4}, RF phase fluctuations of 0.1 degrees, and timing errors in the generation of electrons of 300 fs, the rms fluctuations in beam energy range from 200 keV to 47 keV depending on whether and what type of feedback systems are used. Timing jitter can likewise be reduced by feedback from 54 fs to 15 fs, and jitter in the central current from 47 A to 11 A. Fluctuations in the energy chirp are dominated by the error in the charge per bunch and feedback systems may not be able to stabilize this quantity. Certain other fluctuations, in particular of the energy chirp and timing of the tail of the bunch, remain at similar levels with or without feedback because they are already dominated by shot-to-shot errors in the charge per bunch. Each 1% of fluctuation in bunch charge at time scales too fast to correct by feedback systems contributes roughly 20 keV to energy jitter, 7 fs of timing jitter, and 4.5 A to current jitter. When considering advanced feedback systems, the issue of whether bunch charge or other errors have significant shot-to-shot variations becomes more critical to performance. Reliable estimates of the magnitude and temporal structure of errors will have to be obtained from hardware studies.

Because space charge and coupling between longitudinal and transverse motion have not been considered here, the low-energy portion of the injector should be examined with a more complete physics model to ensure that these effects do not introduce new sensitivies to errors. In particular, it should be confirmed that the effect of changes in the accelerating gradient at the cathode are reasonably consistent with simulations using LiTrack. Expected microstructure in the beam due to effects not considered here, such as coherent synchrotron radiation (CSR) and longitudinal space charge, should also be considered. Another major topic to pursue is the effect of beam jitter and nonuniformity on the FEL performance. This requires both an analysis of the sensitivity of the FEL to jitter in average beam quantities and a series of time-dependent simulations to evaluate the impact of variations across the beam on the delivered photon beam. An initial analysis of sensitivities for an FEL design that uses similar electron beam parameters has already been performed [6].

----

[1] J. Corlett, J. Byrd, W.M. Fawley, M. Gullans, et al., "A high repetition rate VUV-soft X-ray FEL concept", paper TUPMN109 in Proceedings of the Particle Accelerator Conference 2007 (PAC07), June 2007, Albuquerque, NM, IEEE Cat No. 07CH37866 (2007) pp. 1167–69.

[2] K.L.F. Bane and P. Emma, "Litrack: A Fast Longitudinal Phase Space Tracking Code with Graphical User Interface," Proceedings of the 2005 Particle Accelerator Conference (PAC 2005), 16–20 May 2005, pp. 4266–4268.

[3] MATLAB 7.4.0.287 (R2007a), The MathWorks, Inc.

[4] Y. Kim, K. Flöttmann, and T. Limberg, "S2E simulations on jitter for European XFEL project," Proceedings of EPAC 2004, Lucerne, Switzerland, pp. 336–338.

[5] P.N. Ostroumov, B. Mustapha , and V.N. Aseev, "Beam dynamics studies of the 8-GeV superconducting H^- linac," Proceedings of Linac 2006, Knoxville, Tennessee, USA, pp. 420–422.

[6] R. Bartolini, C. Christou, J.-H. Han, I.P.S. Martin et al., "A 1keV FEL driven by a superconducting linac as a candidate for the UK New Light Source," Proceedings of PAC09, Vancouver, BC, Canada, paper TU5RFP062.

[7] A.S. Hedayat, N.J.A. Sloane, and J. Stufken, Orthogonal Arrays, Springer, New York (1999).

[8] M.D. McKay, W.J. Conover and R.J. Beckman, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," Technometrics **21** (1979) pp. 239-245.

[9] G.J. Portmann and A. Zholents, "Ground Motion Measurements at the Advanced Light Source," LBNL ALS report LSAP-226.

[10] B.A. Baklakov, P.K. Lebedev, V.V. Parkhomchuk, A. Sery, A.I. Sleptsov, and V.D. Shiltsev, "Study of seismic vibrations for the VLEPP linear collider," Technical Physics **38** (1993) pp. 894–898.

[11] A. Sery and O. Napoly, "Influence of ground motion on the time evolution of beams in linear colliders," Physical Review E **53** (1996) 5323–5337.

Disclaimer